IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application

Applicant(s): G. Grabarnik et al.
Docket No.: YOR920010748US1
Serial No.: 09/976,543
Filing Date: October 12, 2001
Group: 2142
Examiner: Kelvin Y. Lin

Title: Systems and Methods for Validation, Completion and Construction of Event Relationship Networks

---

APPEAL BRIEF

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Applicants (hereinafter referred to as "Appellants") hereby appeal the final rejection of claims 1-29 of the above-identified application.

REAL PARTY IN INTEREST

The present application is assigned to International Business Machines Corporation, as evidenced by an assignment recorded January 10, 2002 in the U.S. Patent and Trademark Office at Reel 12485, Frame 415. The assignee, International Business Machines Corporation, is the real party in interest.

RELATED APPEALS AND INTERFERENCES

There are no known related appeals or interferences.

STATUS OF CLAIMS

Claims 1-29 stand finally rejected under 35 U.S.C. §103(a). Claims 1-29 are appealed.

STATUS OF AMENDMENTS

There has been no amendment filed subsequent to the final rejection.

SUMMARY OF CLAIMED SUBJECT MATTER

Independent claim 1 recites a computer-based method for use in accordance with an event management system, the method comprising the steps of automatically generating one or more event relationship networks from event data, wherein an event relationship network comprises a graphical representation wherein nodes represent events and links connect correlated nodes; and utilizing the one or more generated event relationship networks to construct one or more correlation rules for use by a correlation engine in the event management system. The present specification provides an illustrative embodiment of the elements of claim 1 at page 8, line 11 through page 9, line 3.

The present specification explains, by way of example at page 7, line 6-10, that the approach taken by the present invention is to describe correlation logic uses a conceptual framework called event relationship networks or ERNs. In one embodiment, an ERN is a directed cyclic graph. Nodes are events and are labeled with the role of the event within the case. Arcs or links from one event to the next indicate that the latter is associated with or correlated with the former. (Specification, page 7, lines 6-10).

More particularly, FIG. 2 illustrates an operational model associated with an event correlation engine according to an embodiment of the present invention. A correlation server 200 comprises a correlation engine 202, an event cache 204, an event throttling module 206, an event repository 208 and a problem reporting system 210. Raw events 212 are received by the correlation server 200. The raw events are preprocessed, via event throttling module 206, such that redundant events are removed. The preprocessed events 214 are then stored in event cache 204. Thus, at any moment, the event cache 204 contains events received during the last period of a predefined duration. The rule-based triggering system (i.e., in accordance with the correlation engine 202 and the correlation rules implemented thereby) examines the content of the event cache 204 and determines whether any trigger rule should fire. The firing of a trigger rule results in the generation of a trouble ticket 218 which is sent on to the problem reporting system 210 for action to be taken

by an operator and/or some response system in the network. Events 216 may be stored for further use in the event repository 208. (Specification, page 8, line 11 through page 9, line 3).

Independent claim 14 is an apparatus claim having similar elements as the above-described claim 1. The apparatus comprises a processor and memory arrangement. FIG. 12 shows an illustrative hardware implementation of a computing system in accordance with which one or more functional components/modules of an event management system may be implemented. As shown, the computer system may be implemented in accordance with a processor 1202, a memory 1204 and I/O devices 1206. (Specification, page 18, lines 5-19).

Independent claim 27 is an article of manufacture claim having similar elements as the above-described claim 1. The article of manufacture comprises a machine readable medium. The present specification provides an illustrative embodiment of the elements of claim 27 at page 18, lines 20-24. More particularly, as explained therein, software components including instructions or code for performing the methodologies described herein may be stored in one or more of the associated memory devices (e.g., ROM, fixed or removable memory) as an article of manufacture and, when ready to be utilized, loaded in part or in whole (e.g., into RAM) and executed by a CPU. (Specification, page 18, lines 20-24).

GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

(I) Whether claims 1, 3, 14, 16 and 27 are unpatentable under 35 U.S.C. §103(a) over D. Mishra, "SNOOP: An Event Specification Language for Active Database System," Thesis from University of Florida, 1991 (hereinafter "Mishra") in view of U.S. Patent No. 6,006,213 issued to Yoshida (hereinafter "Yoshida").

(II) Whether claims 2 and 15 are unpatentable under 35 U.S.C. §103(a) over Mishra in view of Yoshida in further view of U.S. Patent No. 5,345,380 issued to Babson et al. (hereinafter "Babson").

(III) Whether claims 4-7, 11, 13, 17-20, 24 and 26 are unpatentable under 35 U.S.C. §103(a) over Mishra in view of Yoshida in further view of U.S. Patent No. 6,249,755 issued to Yemini et al. (hereinafter "Yemini").

(IV) Whether claims 8-10, 12, 21-23 and 25 are unpatentable under 35 U.S.C. §103(a) over Mishra in view of Yoshida in view of Yemini in further view of Bettini et al., "Testing Complex Temporal Relationship Involving Multiple Granularities and Its Application to Data Mining," ACM 1996 (hereinafter "Bettini").

(V) Whether claims 28 and 29 are unpatentable under 35 U.S.C. §103(a) over Mishra in view of U.S. Patent No. 6,108,698 issued to Tenev (hereinafter "Tenev").

ARGUMENT

Appellants incorporate by reference herein the disclosure of their previous response filed in the present application, namely, the response dated August 31, 2006.

(I) Whether claims 1, 3, 14, 16 and 27 are unpatentable under 35 U.S.C. §103(a) over D. Mishra, "SNOOP: An Event Specification Language for Active Database System," Thesis from University of Florida, 1991 (hereinafter "Mishra") in view of U.S. Patent No. 6,006,213 issued to Yoshida (hereinafter "Yoshida").

Regarding the §103(a) rejections, Appellants assert that the various references, alone or in combination, fail to teach or suggest all of the limitations of claim 1-29, as will be explained below. Furthermore, with regard to the combinations of the various references, Appellants assert that such combinations are improper, as will be explained below.

The Examiner cites Mishra in combination with Yoshida in rejecting independent claims 1, 14 and 27. More particularly, the Examiner cites portions of Mishra as disclosing certain limitations of the independent claims, and cites portions of Yoshida as disclosing certain other limitations of the independent claims. Below, Appellants explain how such portions of Mishra and Yoshida fail to teach or suggest what the Examiner contends that they teach or suggest. While Appellants may refer from time to time to each reference alone in describing its deficiencies, it is to be understood that such arguments are intended to point out the overall deficiency of the cited combination.

The Mishra/Yoshida combination does not disclose automatically generating one or more event relationship networks from event data, wherein an event relationship network comprises a

graphical representation wherein nodes represent events and links connect correlated nodes, as in the claimed invention.

The final Office Action cites the "event compiler" and "event graph" on pages 57 and 58 of Mishra in support of rejecting the claimed step of automatically generating one or more event relationship networks from event data, wherein an event relationship network comprises a graphical representation wherein nodes represent events and links connect correlated nodes. However, the event graph of Mishra is not an event relationship network that can be used to construct one or more correlation rules for use by a correlation engine in an event management system, as in the claimed invention. In fact, as explained at page 57 of Mishra, a rule is actually the input for the graph building algorithm. Then, as explained at page 58, in the event detection technique of Mishra, it is assumed that a event graph already exits and Mishra is detecting instances of the event graph. Thus, this is significantly different than what is claimed.

The Examiner explains in the "Response to Arguments" section of the final Office Action (page 2) that Mishra teaches taking an original rule and creating a "Rule-b" using a constructed tree. Even if this were accurate (which Applicants assert is not the case), such a process does not disclose the steps of the claimed invention. Again, the claimed invention generates one or more event relationship networks from event data, and utilizes the one or more generated event relationship networks to construct one or more correlation rules. By any interpretation of Mishra, this is not what the "event compiler" and "event graph" of Mishra do.

In the Advisory Action, the Examiner again states that Mishra teaches how to construct an event graph, and therefore teaches generating one or more event relationship network from event data. Again, as noted above, the event graph of Mishra is not an event relationship network that can be used to construct one or more correlation rules for use by a correlation engine in an event management system, as in the claimed invention.

The Yoshida reference fails to supplement the above-noted deficiencies of Mishra as applied to claim 1. Accordingly, it is believed that the combined teachings of Mishra and Yoshida fail to meet the limitations of claim 1.

Also, the Examiner has failed to identify a cogent motivation for combining Mishra and Yoshida in the manner proposed. The Examiner provides the following statement of motivation beginning at page 4, first full paragraph of the Office Action:

> It would have been obvious to one ordinary skilled in the art at the time the invention was made to include the teaching of Yoshida for correlation rule used by the correlation engine.
>
> The motivation would be for Mishra to combine with Yoshida is to achieve the extracting a set of patterns frequently appearing in the graph, evaluating the extracted pattern based on the resulting graph size reduction, and outputting a pattern having a good evaluation result (Yoshida, col. 1, l. 60-65).

The Federal Circuit has stated that when patentability turns on the question of obviousness, the obviousness determination "must be based on objective evidence of record" and that "this precedent has been reinforced in myriad decisions, and cannot be dispensed with." In re Sang-Su Lee, 277 F.3d 1338, 1343 (Fed. Cir. 2002). Moreover, the Federal Circuit has stated that "conclusory statements" by an examiner fail to adequately address the factual question of motivation, which is material to patentability and cannot be resolved "on subjective belief and unknown authority." Id. at 1343-1344. There has been no showing in the present §103(a) rejection of claim 1 of objective evidence of record that would motivate one skilled in the art to combine Mishra and Yoshida to produce the particular limitations in question. The above-quoted statement of motivation provided by the Examiner appears to be a conclusory statement of the type ruled insufficient in the In re Sang-Su Lee case. Accordingly, the proposed combination appears to be based primarily on hindsight, with the Examiner attempting to reconstruct the claimed arrangement from disparate references.

For at least these reasons, Appellants assert that independent claim 1 is patentable over the Mishra/Yoshida combination. Independent claims 14 and 27 include limitations similar to those of claim 1, and are therefore believed patentable for reasons similar to those described above with reference to claim 1. Furthermore, Appellants assert that the claims which depend from claim 1 are patentable over the Mishra/Yoshida combination not only for the reasons given above with respect to claim 1, but also because such dependent claims recite patentable subject matter in their own right, as will be set out below.

Regarding claims 3 and 16, it is not clear how "Read rule_definition" in Mishra teaches or suggests obtaining one or more previously generated event relationship networks. (Mishra, page 57, line 10). As noted above, event relationship networks are graphical representations of how events are correlated. At page 57 of Mishra, a rule is the input for the graph building algorithm, and not an event relationship network. It is also not clear how the fact that Mishra describes "Define trees (i.e. nodes) corresponding to rule_event1 and rule_event2 in the forest" teaches or suggests validating the one or more previously generated event relationship networks by removing any nodes or links included therein that are incorrect for a particular application context. Furthermore, it is not clear how "build_tree" of Mishra discloses the step of completing the one or more previously generated event relationship networks by adding any nodes or links thereto that are missing for the particular application context. Lastly, "Create Rule_b" and "merge it in the event_forest" of Mishra do not teach or suggest outputting the one or more validated and completed event relationship networks as the one or more event relationship networks used to construct the one or more correlation rules.

(II) Whether claims 2 and 15 are unpatentable under 35 U.S.C. §103(a) over Mishra in view of Yoshida in further view of U.S. Patent No. 5,345,380 issued to Babson et al. (hereinafter "Babson").

Appellants assert that claims 2 and 15, which depends from claims 1 and 14 respectively, are patentable over the Mishra/Yoshida/Babson combination not only for the reasons given above with respect to claims 1 and 14, but also because such dependent claims recite patentable subject matter in their own right.

Babson does not disclose subjecting the one or more generated event relationship networks to human review prior to utilizing the one or more generated event relationship networks to construct the one or more correlation rules. Instead, the relied-upon portions of Babson refers to "presenting the customer with a plurality of types of nodes, the nodes indicating the determinations and actions allowable for the procedure; receiving from the customer indications of desired

relationships between the desired nodes ...," which is not relevant to the rejection of the claimed features in claims 2 and 15.

It is also asserted that the motivation set forth by the Examiner to combine Babson with Mishra and Yoshida is insufficient under the In re Sang-Su Lee decision (cited above).

(III) Whether claims 4-7, 11, 13, 17-20, 24 and 26 are unpatentable under 35 U.S.C. §103(a) over Mishra in view of Yoshida in further view of U.S. Patent No. 6,249,755 issued to Yemini et al. (hereinafter "Yemini").

Appellants assert that claims 4-7, 11, 13, 17-20, 24 and 26, which depend from claims 1 and 14 respectively, are patentable over the Mishra/Yoshida/Yemini combination not only for the reasons given above with respect to claims 1 and 14, but also because such dependent claims recite patentable subject matter in their own right.

Claims 4 and 17 recite that the validating and completing steps utilize a statistical correlation analysis. The cited portions of Yemini refer to generating a well-formed correlation matrix. However, they do not teach or suggest the claimed features of claims 4 and 17.

Claims 5 and 18 recite that the statistical correlation analysis utilizes pairwise correlation analysis, wherein correlation between a pair of events is measured in accordance with one or more statistical measurements. It is not clear how the correlation matrix including pairs of the form {Pr,t) where Pr is a probability indication, as disclosed in Yemini, teaches or suggests the claimed features of claims 5 and 18.

Regarding claims 6 and 19, it is not clear how applying a filter to remove weakly correlated data after correlations among events are stored in data file teaches or suggests for a particular event relationship network, determining that links in the event relationship network have a confidence level not less than a given threshold.

Claims 7 and 20 recite "splitting the event relationship into correlation paths; for every correlation path, removing a node that has the least number of correlated nodes associated therewith...; and merging correlation paths into one or more event relationship networks..." It is not

clear how the features specifically defined in claims 7 and 20 could be taught or suggested by the cited portions of Yemini.

It is also asserted that the motivation set forth by the Examiner to combine Yemini with Mishra and Yoshida is insufficient under the In re Sang-Su Lee decision (cited above).

(IV) Whether claims 8-10, 12, 21-23 and 25 are unpatentable under 35 U.S.C. §103(a) over Mishra in view of Yoshida in view of Yemini in further view of Bettini et al., "Testing Complex Temporal Relationship Involving Multiple Granularities and Its Application to Data Mining," ACM 1996 (hereinafter "Bettini").

Appellants assert that claims 8-10, 12, 21-23 and 25, which depend from claims 1 and 14 respectively, are patentable over the Mishra/Yoshida/Yemini/Bettini combination not only for the reasons given above with respect to claims 1 and 14, but also because such dependent claims recite patentable subject matter in their own right.

Regarding claims 8 and 21, it is not clear how Bettini at page 73, Figure 2, discloses "utilizing the mined patterns to construct the one or more event relationship networks." Bettini at page 73, Figure 2 illustrates a TAG "which are essentially standard finite automata with the modification that transitions are conditioned not only by input symbols, but also by the values of the associated clock," not one or more event relationship networks. (Bettini, page 68, column 2, line 40 through page 69, column 1, line 2).

With regard to claims 9 and 22, it is not clear how the recited portions of Bettini teach or suggest the claimed features of utilizing a statistical correlation analysis to mine patterns. The relationship between X_1 and X_0 dictating that the event assigned to X_0 must happen during the first month of a year does not disclose the claimed features of claims 9 and 22.

Regarding claims 10 and 23, Bettini at page 71, Figure 1 does not teach or suggest of utilizing pairwise correlation analysis, wherein correlation between a pair of events is measured in accordance with one or more statistical measurements. Figure 1 of Bettini only shows two event structures, and no where does Figure 1 show the correlation between a pair of events measured in accordance with one or more statistical measurements.

It is also asserted that the motivation set forth by the Examiner to combine Bettini with Mishra, Yoshida and Yemini is insufficient under the In re Sang-Su Lee decision (cited above).

(V) Whether claims 28 and 29 are unpatentable under 35 U.S.C. §103(a) over Mishra in view of U.S. Patent No. 6,108,698 issued to Tenev (hereinafter "Tenev").

Appellants assert that claims 28 and 29, which depend from claim 1, are patentable over the Mishra/Tenev combination not only for the reasons given above with respect to claim 1, but also because such dependent claims recite patentable subject matter in their own right.

Appellants added, in the response dated February 10, 2006, new claims 28 and 29. In the final Office Action, the Examiner introduces the Tenev reference in combination with Mishra to reject new claims 28 and 29. In particular, the final Office Action points to column 9, lines 32-45, of Tenev to reject claim 28 and column 13, lines 28-39, of Tenev to reject claim 29.

Column 9, lines 32-45, of Tenev read:

> FIG. 6 illustrates features of directed graph data structure 330 that are relevant to the operations performed by grapher routines 320 in relation to the expansion flags.
>
> Identifier (ID) mapping structure 350 maps from element IDs to pointers. The element IDs include node IDs and link IDs. Structure 350 makes it possible for every node and link in memory to be specified by an ID which can be validated in constant time and nearly always created in constant time; structure 350 avoids the need to use pointers except within directed graph data structure 330. Although implemented as two arrays of pointers, one indexed by the node IDs and the other by link IDs, structure 350 could also be, for example, a lookup table in which each entry includes an ID and a pointer.

No where does this portion of Tenev or any other portion of Tenev teach or suggest computing a first correlation metric and a second correlation metric, the second correlation metric being representative of a correlation between events that is stronger than a correlation between events represented by the first correlation metric, as in claim 28. In fact, Appellants fail to see any discussion whatsoever in Tenev regarding such correlation metrics.

Further, column 13, lines 28-39, of Tenev read:

> FIG. 9 shows a sequence of representations of the graph shown in box 202 in FIG. 4 that could be presented as a result of operations like those described above in relation to FIG. 7, without creating or removing any nodes. Each representation could result from a respective iteration in FIG. 7, either an iteration through boxes 364 through 382 in response to an expand signal or an iteration through boxes 390 and 392 in response to a contract signal. In each representation, a node feature with a "+" indicates that the represented node is contracted in the tree, while a node feature with a "-" indicates that the represented node is expanded in the tree.

No where does this portion of Tenev or any other portion of Tenev teach or suggest specifying an event data window within which event data is considered, as in claim 29. The "+" and "-" mentioned above in Tenev relate to the graph itself and not to any event data or an event data window within which event data is considered, as claimed.

In the Advisory Action, the Examiner points to column 6, line 51 through column 7, line 6 of Tenev as disclosing the limitations of claim 28. However, these cited portions of Tenev also fail to teach or suggest computing a first correlation metric and a second correlation metric, the second correlation metric being representative of a correlation between events that is stronger than a correlation between events represented by the first correlation metric, as recited in claim 28. With regard to claim 29, the Examiner again refers to column 9, lines 32-45 and FIG. 6 of Tenev as disclosing the limitations of claim 29. As noted above, no where does this portion of Tenev or any portion of Tenev teach or suggest the limitations as recited in claim 29.

It is also asserted that the motivation set forth by the Examiner to combine Tenev with Mishra is insufficient under the In re Sang-Su Lee decision (cited above).

In view of the above, Applicants believe that claims 1-29 are in condition for allowance, and respectfully request withdrawal of the §103(a) rejections.

Respectfully submitted,

Date: January 9, 2007

William E. Lewis
Attorney for Applicant(s)
Reg. No. 39,274
Ryan, Mason & Lewis, LLP
90 Forest Avenue
Locust Valley, NY 11560
(516) 759-2946

APPENDIX

1. A computer-based method for use in accordance with an event management system, the method comprising the steps of:

automatically generating one or more event relationship networks from event data, wherein an event relationship network comprises a graphical representation wherein nodes represent events and links connect correlated nodes; and

utilizing the one or more generated event relationship networks to construct one or more correlation rules for use by a correlation engine in the event management system.

2. The method of claim 1, further comprising the step of subjecting the one or more generated event relationship networks to human review prior to utilizing the one or more generated event relationship networks to construct the one or more correlation rules.

3. The method of claim 1, wherein, when one or more previously generated event relationship networks are available, the step of automatically generating one or more event relationship networks comprises:

obtaining one or more previously generated event relationship networks;

validating the one or more previously generated event relationship networks by removing any nodes or links included therein that are incorrect for a particular application context;

completing the one or more previously generated event relationship networks by adding any nodes or links thereto that are missing for the particular application context;

outputting the one or more validated and completed event relationship networks as the one or more event relationship networks used to construct the one or more correlation rules.

4. The method of claim 3, wherein the validating and completing steps utilize a statistical correlation analysis.

5. The method of claim 4, wherein the statistical correlation analysis utilizes pairwise correlation analysis, wherein correlation between a pair of events is measured in accordance with one or more statistical measurements.

6. The method of claim 3, wherein the validating step comprises, for a particular event relationship network, determining that links in the event relationship network have a confidence level not less than a given threshold.

7. The method of claim 3, wherein the validating step, for a particular event relationship network, comprises:

splitting the event relationship network into correlation paths;

for every correlation path, removing a node that has the least number of correlated nodes associated therewith until every node is fully correlated with every other node; and

merging correlation paths into one or more event relationship networks such that every path in a resulting event relationship network has every node fully correlated with every other node in the path.

8. The method of claim 1, wherein, when one or more previously generated event relationship networks are not available, the step of automatically generating one or more event relationship networks comprises:

mining patterns from the event data;

utilizing the mined patterns to construct the one or more event relationship networks;

outputting the one or more event relationship networks constructed from the mined patterns as the one or more event relationship networks used to construct the one or more correlation rules.

9. The method of claim 8, wherein the constructing step utilizes a statistical correlation analysis to mine patterns.

10. The method of claim 8, wherein the statistical correlation analysis utilizes pairwise correlation analysis, wherein correlation between a pair of events is measured in accordance with one or more statistical measurements.

11. The method of claim 1, wherein the event data is obtained from an event log representing historical events associated with a particular system being managed by the event management system.

12. The method of claim 1, wherein the one or more event relationship networks comprise annotations relating to statistical correlation between nodes.

13. The method of claim 1, wherein the event data is preprocessed prior to use in generating the one or more event relationship networks by removing at least a portion of any redundant events.

14. Apparatus use in accordance with an event management system, the apparatus comprising:

at least one processor operative to: (i) automatically generate one or more event relationship networks from event data, wherein an event relationship network comprises a graphical representation wherein nodes represent events and links connect correlated nodes; and (ii) utilize the one or more generated event relationship networks to construct one or more correlation rules for use by a correlation engine in the event management system; and

memory, coupled to the at least one processor, which stores at least one of the event data and the one or more event relationship networks.

15. The apparatus of claim 14, wherein the at least one processor is further operative to permit the operation of subjecting the one or more generated event relationship networks to human review prior to utilizing the one or more generated event relationship networks to construct the one or more correlation rules.

16. The apparatus of claim 14, wherein, when one or more previously generated event relationship networks are available, the operation of automatically generating one or more event relationship networks comprises:

obtaining one or more previously generated event relationship networks;

validating the one or more previously generated event relationship networks by removing any nodes or links included therein that are incorrect for a particular application context;

completing the one or more previously generated event relationship networks by adding any nodes or links thereto that are missing for the particular application context;

outputting the one or more validated and completed event relationship networks as the one or more event relationship networks used to construct the one or more correlation rules.

17. The apparatus of claim 16, wherein the validating and completing operations utilize a statistical correlation analysis.

18. The apparatus of claim 17, wherein the statistical correlation analysis utilizes pairwise correlation analysis, wherein correlation between a pair of events is measured in accordance with one or more statistical measurements.

19. The apparatus of claim 16, wherein the validating operation comprises, for a particular event relationship network, determining that links in the event relationship network have a confidence level not less than a given threshold.

20. The apparatus of claim 16, wherein the validating operation, for a particular event relationship network, comprises:

splitting the event relationship network into correlation paths;

for every correlation path, removing a node that has the least number of correlated nodes associated therewith until every node is fully correlated with every other node; and

merging correlation paths into one or more event relationship networks such that every path in a resulting event relationship network has every node fully correlated with every other node in the path.

21. The apparatus of claim 14, wherein, when one or more previously generated event relationship networks are not available, the step of automatically generating one or more event relationship networks comprises:

mining patterns from the event data;

utilizing the mined patterns to construct the one or more event relationship networks;

outputting the one or more event relationship networks constructed from the mined patterns as the one or more event relationship networks used to construct the one or more correlation rules.

22. The apparatus of claim 21, wherein the constructing operation utilizes a statistical correlation analysis to mine patterns.

23. The apparatus of claim 21, wherein the statistical correlation analysis utilizes pairwise correlation analysis, wherein correlation between a pair of events is measured in accordance with one or more statistical measurements.

24. The apparatus of claim 14, wherein the event data is obtained from an event log representing historical events associated with a particular system being managed by the event management system.

25. The apparatus of claim 14, wherein the one or more event relationship networks comprise annotations relating to statistical correlation between nodes.

26. The apparatus of claim 14, wherein the event data is preprocessed prior to use in generating the one or more event relationship networks by removing at least a portion of any redundant events.

27. An article of manufacture for use in accordance with an event management system, the article comprising a machine readable medium containing one or more programs which when executed implement the steps of:

automatically generating one or more event relationship networks from event data, wherein an event relationship network comprises a graphical representation wherein nodes represent events and links connect correlated nodes; and

utilizing the one or more generated event relationship networks to construct one or more correlation rules for use by a correlation engine in the event management system.

28. The method of claim 1, further wherein automated generation of at least one of the one or more event relationship networks comprises use of an automated pairwise statistical correlation procedure which is configured to compute a first correlation metric and a second correlation metric, the second correlation metric being representative of a correlation between events that is stronger than a correlation between events represented by the first correlation metric.

29. The method of claim 1, further wherein automated generation of at least one of the one or more event relationship networks comprises specifying an event data window within which event data is considered.

EVIDENCE APPENDIX

None.

RELATED PROCEEDINGS APPENDIX

None.